

10/718,009  
MO-892

# (12) UK Patent Application (19) GB (11) 2 279 165 (13) A

(43) Date of A Publication 21.12.1994

(21) Application No 9411286.9

(22) Date of Filing 06.08.1994

(30) Priority Data

(31) 08077227 (32) 14.06.1993 (33) US

(51) INT CL<sup>5</sup>  
G06F 15/403

(52) UK CL (Edition M )  
G4A AUDB

(56) Documents Cited  
None

(58) Field of Search  
NO SEARCH POSSIBLE

(71) Applicant(s)

Hewlett-Packard Company

(Incorporated in USA - California)

3000 Hanover Street, Palo Alto, California 94304,  
United States of America

(74) Agent and/or Address for Service  
Williams, Powell & Associates  
34 Tavistock Street, LONDON, WC2E 7PB,  
United Kingdom

(72) Inventor(s)

Surajit Chaudhuri

Kyuseok Shim

(54) Query optimization in a relational database system.

(57) Database applications typically need to invoke foreign functions or to access data that is not stored in the database. The invention provides a comprehensive approach to cost-based optimization of relational queries in the presence of such foreign functions. The optimization takes into account semantic information about foreign functions using a declarative rule language (e.g., SQL) to express such semantics. Procedures for applying the rewrite rules and for generating the execution space of equivalent queries are described. Procedures to obtain an optimal plan from this execution space are also described. Moreover, necessary extensions to the cost model that are needed in the presence of foreign functions are described.

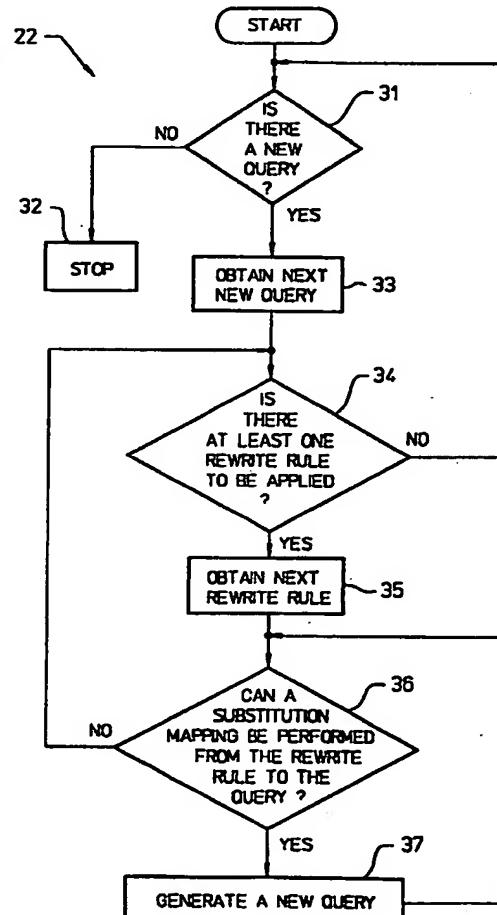


FIG. 3

GB 2 279 165 A

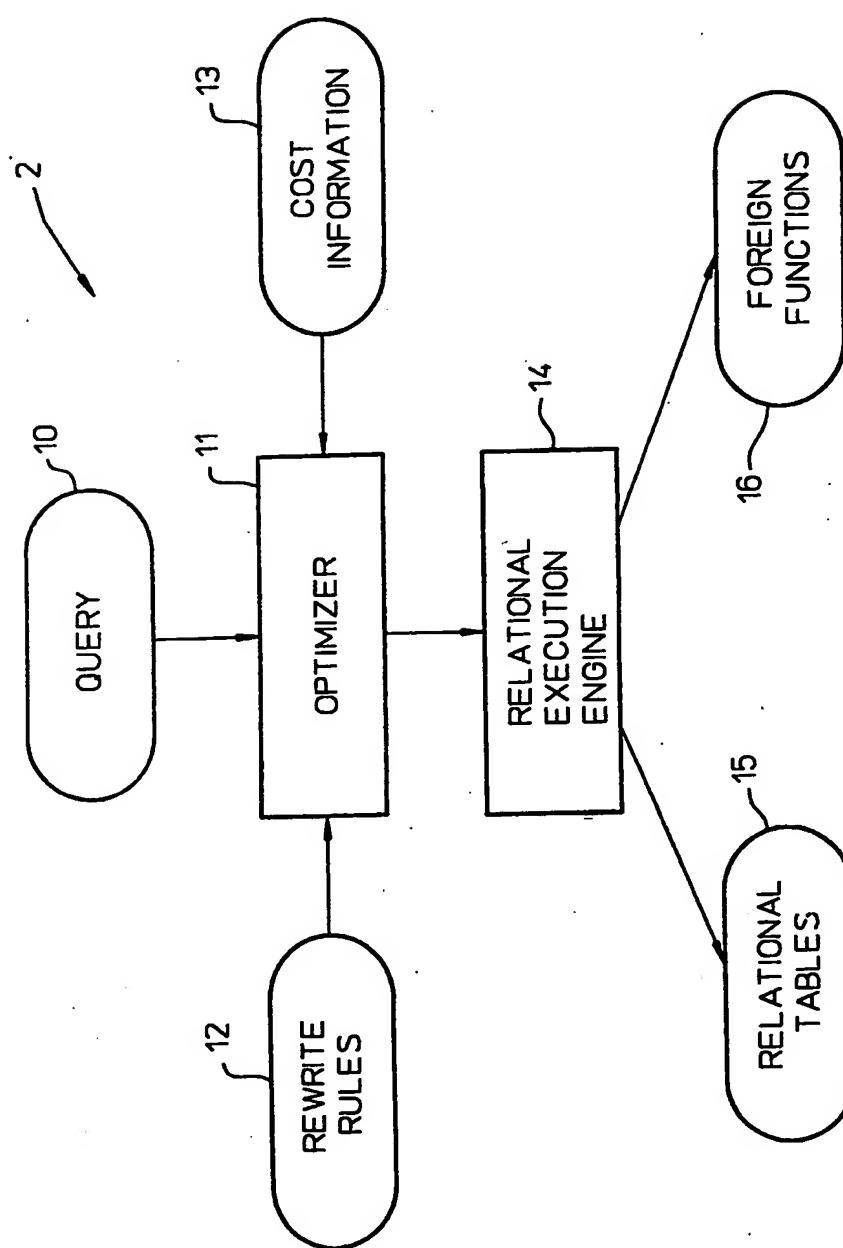


FIG. 1

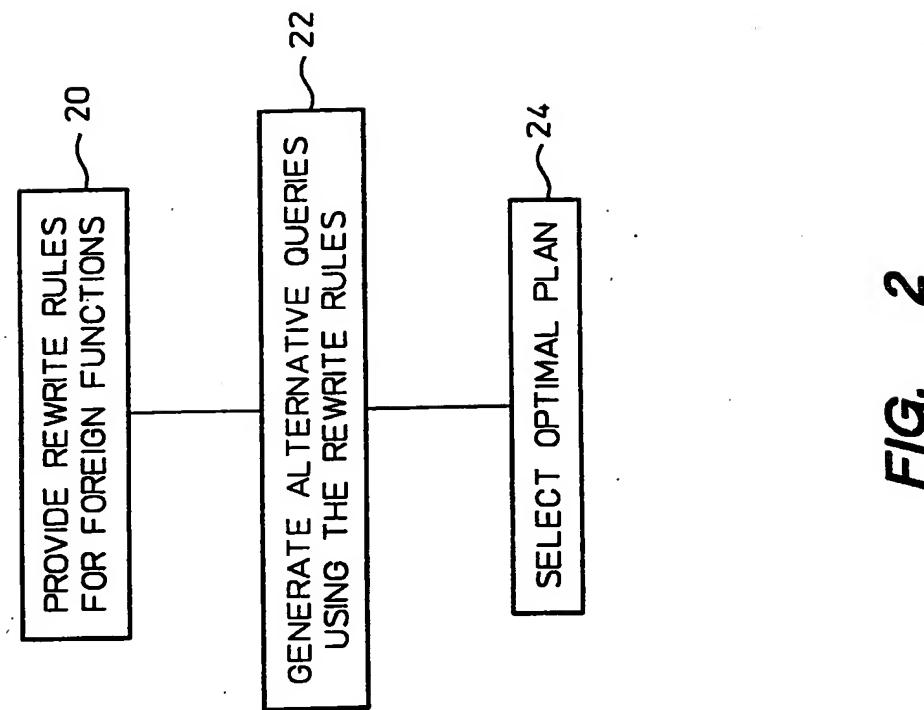
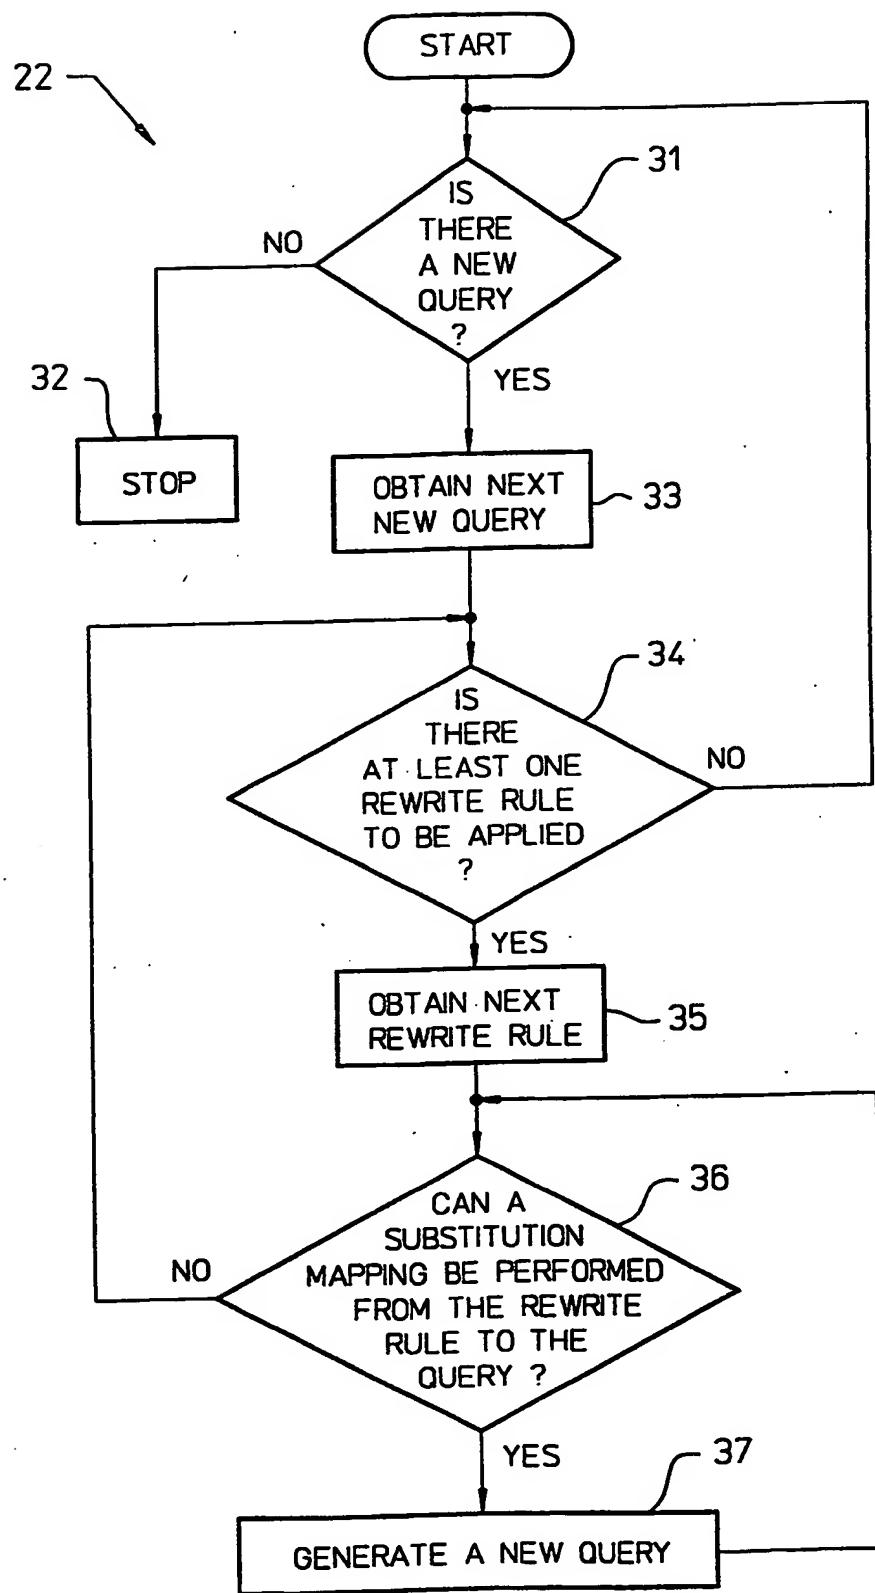


FIG. 2

**FIG. 3**

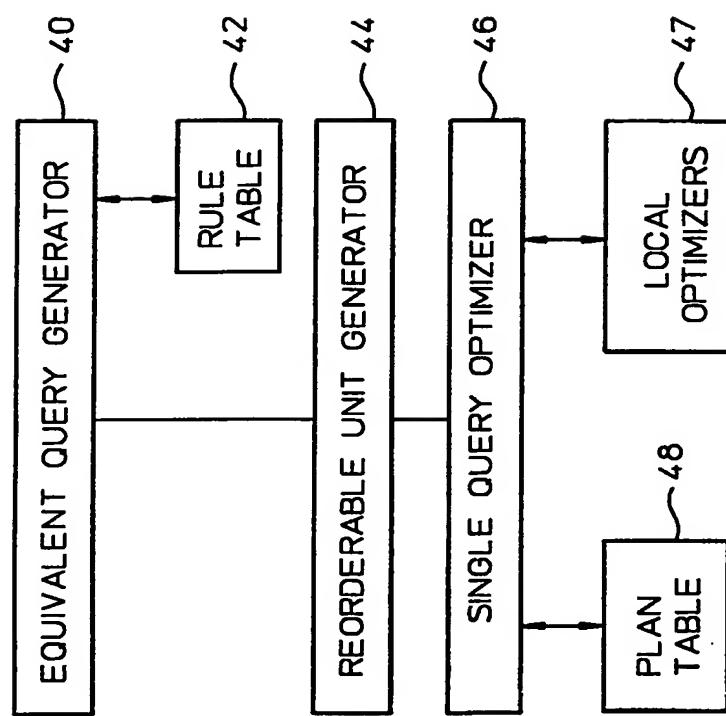


FIG. 4

## QUERY OPTIMIZATION IN A RELATIONAL DATABASE SYSTEM

The present invention relates to a relational database system and, more particularly, to query optimization techniques for a relational database system having foreign functions.

5           Relational database systems provide the ability to conveniently query the data stored in their database. However, in many applications, there is a need to integrate data and operations that are external to the database (referred to as foreign functions). For example, it would be convenient to invoke mathematical functions and UNIX library functions as part of a relational query. Moreover, for  
10          many problem domains, highly tuned applications exist. The ability to exploit such existing applications is important since redevelopment can be prohibitively expensive. Also, for many applications, only part of the data that is needed may be stored in the database. Additional data may reside externally. Access to external data is provided by a set of interface routines.

15          As an example, many specialized Geographic Information Systems (GIS) are available today that provide the ability to store and access geographic data. On the other hand, information on attributes (e.g., population of a city) is usually stored in a relational database. Therefore, for GIS applications, the ability to use a relational query language as well as the ability to invoke functions provided by  
20          the GIS package is important. In general, the ability to invoke foreign functions in a relational query is important to develop applications.

To illustrate the key challenges to optimization introduced by foreign functions, consider two examples taken from an earlier application described in Kolovson et al., *Interoperability of spatial and attribute data managers: A case*

study, Proceedings of the 3rd International Symposium on Large Spatial Databases, Miami, Florida, 1992. This known application was built in the Papyrus project, see Connors et al., *The papyrus integrated data server*, Proceedings of the First International Conference on Parallel and Distributed Systems, Miami, Florida, 1991. The application allows one to access information about businesses and their locations in the Bay Area of northern California. The application is built on top of the ETAK MapEngine and a relational storage manager. ETAK Inc. of Menlo Park, California is a company that designs vehicle navigation equipment and produces digital map databases.

10       The MapEngine is a geographic data manager that provides the ability to store and query maps. The MapEngine stores the locations of the business establishments in the Bay Area in a file Map. The relational database is used to store attribute information about businesses (e.g., type of business) in a table Business. Each tuple of this table also stores an additional attribute which acts as a key for the MapEngine. The latter uses this key to look up the location of the business. Similarly, each record in MapEngine points to the tuple in the table Business of the relational database where the attribute information about the corresponding business establishment is maintained. Thus, queries are permitted to span the relational system as well as the MapEngine.

15       Example 1: The aim of this example is to highlight the importance of the semantic information associated with the foreign functions, for optimization. The MapEngine provides a function to access all points (a point or a window is represented as a single argument, although a variety of representations is possible) in a map (Map) and a boolean function (Inside) to test whether a given point is within a window. The function Inside is an arithmetic function. The MapEngine also provides an additional function Mapclip that, given a window, returns all points in the map that are in that window. Consider a query to find, given a window, all points in the map that are in the window. The query can be answered by invoking Map and testing that each of the retrieved locations is inside the window

(by using the function Inside). However, the use of the fact that the query can be answered by invoking Mapclip is significant since the use of Mapclip can greatly reduce the cost of evaluation of the query.

5        Example 2: The aim of this example is to highlight that the decision to modify a given query by semantic optimization may need to be cost-based. Consider the problem of finding all restaurants in downtown Palo Alto. This query can be answered by selecting all restaurants from the table Business and then performing a Mapclip.

10      However, MapEngine also has a file Map\_Restaurant consisting of all the establishments that are restaurants. Therefore, this semantic information can be used to answer the queries. One could invoke Map\_Restaurant to obtain all restaurants in the Bay Area and then can select those in downtown Palo Alto by invoking Inside. These 15      two queries are equivalent, but the optimal plan for one of the queries may be better, even by an order of magnitude, compared to the optimum plan for the other query, depending on whether the indexing effect of restricting locations to downtown Palo Alto is more effective than indexing based on restricting the businesses to 20      be restaurants.

Example 1 illustrates that in the presence of foreign functions, there may be multiple ways to answer the same query and such semantic information is extremely valuable for query optimization and must be captured. Example 2 illustrates that application of such semantic information for query optimization may 25      need to be cost-based.

The ability to answer relational queries efficiently relies on the repertory of evaluation options and an optimizer to choose among these options. Therefore, when relational queries have the ability to invoke foreign functions, the database system needs to provide sufficient evaluation options and necessary extensions to 30      the optimizer so that queries involving foreign functions can be effectively optimized. Existing optimizers, however, are unable to satisfy this need. Of

course, there are other dimensions to the problem of invoking foreign functions (e.g., format conversion, supporting complex objects), but the focus of the invention is on optimization and related issues only.

In recent years, several extensible systems have been proposed with varying degrees of extensibility. See, e.g., Carey et al., *Extensible database management systems*, ACM-SIGMOD Record, Dec. 1990; Grefe et al., *The exodus optimizer*, Proceedings of the 1987 ACM-SIGMOD Conference on the Management of Data, pp. 160-172, San Francisco, CA, May 1987; Hass et al., *Extensible query processing in starburst*, Proceedings of the 1989 ACM-SIGMOD Conference on the Management of Data, pp. 377-388, Portland, OR, June 1989, Stonebraker et al., *On rules, procedures, caching and views in database systems*, Proceedings of the 1990 ACM-SIGMOD Conference on the Management of Data, pp. 281-290, Atlantic City, NJ, May 1990. Rewrite languages for optimization in extensible systems are also known. See, e.g., Pirahesh et al., *Extensible/rule based query optimization in starburst*, Proceedings of the 1992 ACM-SIGMOD Conference on the Management of Data, pp. 39-48, San Diego, CA, May 1992; Lohman, *Grammer-like functional rules for representing query optimization alternatives*, Proceedings of the 1988 ACM-SIGMOD Conference on the Management of Data, pp. 18-27, Chicago, IL, June 1988; Grefe et al., *supra*.  
The references of Pirahesh et al., Lohman, and Grefe et al. are hereby incorporated by reference.

A major disadvantage of these known rewrite languages is that they lack the ability to operate on foreign functions which appear in a query. Another disadvantage is that the existing rewrite languages resemble an intermediate-level programming language (e.g., C) which is cumbersome and complicated to work with. Yet another disadvantage of the existing rewrite languages is that the optimization techniques are dependent on the rewrite rules, thus complicating optimization.

Query optimization in the presence of foreign function was examined in Chimenti et al., *Towards an open architecture for LDL*, Proceedings of the 15th International VLDB Conference, pp. 195-203, Amsterdam, August 1989, which is

hereby incorporated by reference. In this reference LDL programs were extended to allow foreign tables and to define cost descriptors for the foreign tables. The LDL system did not make use of rewrite rules containing semantic information about the foreign functions.

5        Thus, prior attempts to optimize queries invoking foreign functions have not made use of a high level language for rewrite rules to express semantic relationships of foreign functions. Nor have prior attempts been able to guarantee that an optimal plan can be obtained using rewrite rules and cost models for foreign functions. Therefore, the problem of cost-based optimization of relational  
10      queries in the presence of such foreign functions, has not previously been addressed satisfactorily.

15      An objective of the invention is to provide improved cost-based optimization of relational queries in the presence of foreign functions.

Broadly speaking, the invention concerns an optimization approach that takes into account semantic information about foreign functions.

20      More particularly, the invention is a comprehensive approach for query optimization in the presence of foreign functions. Rewrite rules are used to express semantics of foreign functions. The rewrite rules are specified using an extension to the query language. The rewrite rules present an optimizer with a space of equivalent queries. The invention also ensures that an optimal plan (from the plans for the original query as well as the alternative queries generated by rewrite rules) is obtained.

25      The invention has many other aspects which are completely described in the detailed description. For example, the invention makes use of a rule-independent procedure for application of the rewrite rules to produce equivalent queries, and exploits the commonality among queries when optimizing the alternative queries.

30      The ability to invoke foreign functions in a relational query is important for many applications since it provides them the opportunity to exploit existing code

and data that is external to the database. The present invention enables queries that invoke foreign functions to be efficiently and effectively optimized.

The present invention will be readily understood by the following detailed description of exemplary embodiments of the invention in conjunction with the accompanying drawings, wherein like reference numerals designate like structural elements, and in which:

Figure 1 is a block diagram of a relational database system in accordance with the invention;

Figure 2 is a basic flowchart illustrating the aspects of the invention;

Figure 3 is a flowchart illustrating a closure procedure according to the second aspect of the invention which generates equivalent queries; and

Figure 4 is a block diagram of an embodiment of the optimization procedure according to the invention.

Embodiments of the invention are discussed below with reference to Figures 1-4. However, those skilled in the art will readily appreciate that the detailed description given herein with respect to these figures is for explanatory purposes as the invention extends beyond these limited embodiments.

The invention concerns a relational database management system which includes method and apparatus for optimizing queries that invoke foreign functions. Semantic information is specified in a declarative way (using a simple extension to SQL) by using rewrite rules that express semantics of foreign functions. The incorporation of semantic information implicitly enhances the space of choices available to the optimizer. An optimal plan is then picked in a cost-based fashion that considers all alternate phrasings of the query implied by the semantic information. Additionally, heuristics can be incorporated in the optimization approach.

Figure 1 is a block diagram illustrating a relational database system 2 in accordance with the invention. The relational database system 2 receives a query

10 which is to be optimized. The query is supplied to an optimizer 11 which  
optimizes the query 10 in accordance with rewrite rules 12 and cost information  
13. In particular, the optimizer 11 receives an input query to be processed,  
generates alternative queries using the rewrite rules 12, and then selects an  
5 "optimal" plan from a group of queries including the input query and the  
alternative queries. The relational database system 2 also includes a relational  
execution engine 14 which can access relational database tables 15 and foreign  
functions 16. The rewrite rules 12 are associated with either the foreign functions  
16 or the database tables 15. The specifier of the rewrite rules 12 assures that  
10 over all databases the queries on either side of the rewrite rule are equivalent.  
The foreign functions 16 denote data and operations which are external to the  
database (e.g., foreign conditions, foreign tables or foreign functions).

Figure 2 is a basic flowchart illustrating the aspects of the invention. As  
shown in Figure 2, there are three aspects of the invention. The first aspect  
15 concerns the provision 20 of rewrite rules for foreign functions. The second  
aspect concerns generating 22 alternative queries using the rewrite rules. The  
third aspect concerns selection 24 of an optimal plan using a cost-based approach.  
Each of these aspects is described in detail below.

20 First Aspect of the Invention  
A first aspect of the invention relates to the provision 20 of rewrite rules  
12 for foreign functions 16 in a relational database system 2. A rewrite rule is  
used to rewrite a query including foreign functions (e.g., tables) as an equivalent  
query which has a different format but still includes foreign functions (e.g.,  
tables). The equivalent query may be executed more efficiently. Hence, the  
25 rewrite rules for the foreign functions enhance the optimization potential of query  
optimizers by taking into account semantic information about foreign functions.

The focus of the invention is directed toward conjunctive queries.  
Conjunctive queries correspond to a subset of SQL which has the following form:

30           SELECT columnlist FROM Tablelist  
               WHERE cond1 AND ... AND condn

Observe that the WHERE clause is a conjunction of conditions cond1 ... condn.

Every conjunctive query is a flattened select-project-join (SPJ) query. This subset of SQL is most widely used. Although the embodiments of the invention discussed below concentrate on conjunctive queries, the invention is generally applicable to all SQL queries.

5        For notational convenience, conjunctive queries may be represented as domain calculus expressions as is done in nonrecursive Datalog. See, e.g., Ullman, *Principals of Database and Knowledge-Bases Systems*, Computer Science Press, 1989. In a domain calculus expression, a conjunctive query is represented as a set of conjuncts (literals). Each conjunct is a table name with arguments.

10      There are no explicit equality clauses in such a domain-calculus expression. Instead, the equalities are implicitly represented as equality of variables in the expression. Like SQL, the result of such a domain-calculus expression is a bag of tuples. This approach differs from the approach used in Deductive Databases, where a set semantics is associated with such a notation.

15      No special syntax in the domain calculus notation is needed to refer to foreign functions. A reference to foreign function in the domain calculus notation appears simply as another conjunct. Therefore, foreign functions are modeled as foreign tables (the terms foreign functions and foreign tables are used interchangeably). The following example indicates how in the domain calculus representation, any reference to foreign functions in the query is represented uniformly as a conjunct although the reference may occur as a condition, table or function in the SQL query.

20      25      30

Example 3:

Consider a slightly modified version of the query that was informally stated in Example 1. Assume that a table BUSINESS that has five attributes: NAME, TYPE, EARNING, SIZE and ETAKID. The map on MapEngine is modeled as a foreign table MAP consisting of attributes ETAKID and LOCATION. The attribute ETAKID in both the tables refers to the key in the MapEngine. Recall from Example 1 that Inside acts as a condition that checks whether a point is within a window. Therefore, it can

be represented as a condition in the WHERE clause of the query.

Finally, a foreign function EXPECTED-REVENUE takes the size of a restaurant as an input argument and estimates the average expected earning of a restaurant. The following query finds all restaurants that are in the map in the window w and who are making more profit than the expected amount.

5

10

```

SELECT BUSINESS.NAME, MAP.LOCATION
FROM BUSINESS, MAP
WHERE BUSINESS.TYPE = 'Restaurant'
AND BUSINESS.ETAKID = MAP.ETAKID
AND INSIDE(w, MAP.LOCATION)
AND BUSINESS.EARNING > EXPECTED-
REVENUE(BUSINESS.SIZE)

```

15

The domain-calculus representation for the same query is:

```

Query(name) :- Business(name, "Restaurant", earn, size, eid),
Map(eid, location), Insidebb(w, location), Exp_Revbf(size, exp),
earn > exp

```

20

Observe that depending on whether the foreign function occurs as a table (MAP) or as a function (EXPECTED-REVENUE), its representation in the domain-calculus notation varies. Namely, the foreign function occurring as a table has arguments for each of its attributes, whereas the foreign function corresponding to a function has one argument position for every input argument position and one argument position for every output argument position. The superscripts are used to indicate the safety constraints on the foreign functions.

25

30

For an n-ary conjunct, the superscript is an n-ary list, one for each argument position. The superscript b (bound) indicates whether the argument must be passed a value, before the foreign function is invoked. Otherwise, the superscript is f (free). For example, the conjunct Inside requires both its arguments to be bound. For simplicity, the superscript is omitted below if all the arguments may be free, or where the binding information is not relevant.

In the domain-calculus representation, references to all tables, foreign or local (i.e., stored), appear similar in the query. Even so, the distinction between a foreign table and a database table is important for query evaluation as well as for query optimization.

5       The objective of the rewrite rules is to capture semantic information associated with foreign tables and their relationship to database tables. The representation of rewrite rules is declarative. The declarativeness of the rewrite rules not only makes it possible to provide a formal semantics and but also facilitates rule-independent algorithms for application of rules and generating 22  
10 alternatives to the given query.

The representation of rewrite rules requires only simple extensions to the query language SQL. Roughly speaking, a rewrite rule has the format REWRITE QUERY1 AS QUERY2 where QUERY1 and QUERY2 are relational queries. It is required that the result of the queries have the same arity (i.e., number of  
15 columns). An important point is that application of such a rule, by default, generates a new query which is considered by the optimizer 11 as an alternative to the given query 10 (though one could specify that only the new query is to be considered). In any case, the final selection 24 of a query from the given query  
10 and the alternatives produced 22 by the rewrite rules is made based on  
20 cost-based optimization.

The following notation is used for the rewrite rules.

$$E_l(x, y) \Rightarrow E_r(x, z)$$

The expressions  $E_l(x, y)$  and  $E_r(x, z)$  are conjunctive expressions and will be called the left-hand side (lhs) and the right-hand side (rhs) of the rule respectively.  
25      The variables x, y and z are ordered sets of variables. The set of variables x which occurs in either side of the rewrite rule are called universal variables. As will be discussed in more detail below, the rewrite rule says that the left-hand side of the rule can be replaced by the right-hand side of the rule "over universal variables." In addition, the notation  $\Leftrightarrow$  is used to indicate two rules all at once  
30      (i.e., bidirectional rules).

As a first example of a rewrite rule, consider the rule informally used in Example 1. This rewrite rule can be represented in a domain calculus notation as:

Map(eid, loc), Inside(window,loc)  $\Rightarrow$  Mapclip(eid, loc, window)

Note that the safety constraint for Mapclip is (ffb). In this rewrite rule, the variables eid, loc and window are all universal variables.

5

As a second example of a rewrite rule, the following rule was used informally in Example 2.

Business(name, "Restaurant", earn, size, eid), Map(eid, loc)  $\Leftrightarrow$   
Map\_Restaurant(eid, loc)

10

This rule says that in order to obtain locations of all restaurants, one can either take a join between Business and Map or one can use the ETAK file Map\_Restaurant. Here, eid and loc are universal variables.

15

As a third example of a rewrite rule, the following rule for MapEngine says that instead of separately checking whether a point belongs to two given windows, one can check whether the point belongs to the intersection of windows.

Inside(w1,point), Inside(w2,point)  $\Rightarrow$  Inside(w,point), Intersect(w1,w2,w)

Using this rewrite rule, the problem of finding all businesses in multiple windows can be reduced to the problem of finding all businesses in the intersection of the windows.

20

As a fourth example of a rewrite rule, the following rewrite rule facilitates query optimization. It is often useful to be able to indicate to the optimizer that an index exists so that it can be used when appropriate. For example, assume there is an index on Map for a given eid. A rewrite rule based on this assumption would then be as follows:

25

Map(eid, loc)  $\Leftrightarrow$  Mapwithid<sup>bf</sup>(eid,loc)

The safety constraint on Mapwithid is (bf) which requires the eid to be specified before it is invoked. Just as in traditional optimization, use of an index is cost-based, so must the choice between the given query and the query obtained by an application of the above rule.

30

Generally speaking, the rewrite rule is used to generate alternatives to the given query by applying the rules from left to right. The formal semantics

associated with a rewrite rule has the following two components: equivalence and directionality. Two queries are equivalent if they result in the same bag of tuples over any database.

First, a rewrite rule assures the following equivalence. For a rewrite rule,  
 5       $E_l(x, y) \Rightarrow E_r(x, z)$ , over any database, the queries  $Q_l$  and  $Q_r$ , as defined below, result in the same bag of tuples.

$$Q_l(x) :- E_l(x, y)$$

$$Q_r(x) :- E_r(x, z)$$

By virtue of the above equivalence, a rewrite rule can be used to derive an  
 10     equivalent query by "substituting" subexpressions in a query. Observe that only the universal variables occur as projection variables of  $Q_r$  and  $Q_l$ .

Next, a rewrite rule also specifies directionality, as indicated by the arrow ( $\Rightarrow$ ). The arrow is used to indicate that only an occurrence of the left-hand side of the rule should be "substituted" by the corresponding occurrence of the right-hand  
 15     side (and not vice-versa) to generate equivalent queries.

Consider the first example of a rewrite rule discussed above. The semantics imply that, over any database, queries  $Q_l$  and  $Q_r$  must result in same bag of tuples over any database.

$$Q_l(eid, loc, window) :- Map(eid, loc), Inside(window, loc)$$

$$20 \quad Q_r(eid, loc, window) :- Mapclip(eid, loc, window)$$

This rewrite rule has an occurrence in the query  $Q$ , given below. The left-hand side of the rule matches the second and the third conjuncts of  $Q$ . By replacing those conjuncts with the corresponding substitution for the right-hand side of the rule, the query  $Q'$  is obtained.

$$25 \quad Q(name, loc) :- Business(name, "Restaurant", earn, eid), Map(eid, loc), \\ Inside(w, loc), Intersect(w1, w2, w)$$

$$Q'(name, loc) :- Business(name, "Restaurant", earn, eid), Mapclip(eid, loc, w), \\ Intersect(w1, w2, w)$$

The directionality of the semantics imply that the rule in the third example of the  
 30     rewrite rules cannot be applied to query  $Q$  although an application will result in an equivalent query.

In the preferred embodiment of the first aspect of the invention, the rewrite rules are provided 20 to the relational database system 2 using a high level language such as an extension of the query language SQL. Thus, although the domain calculus notation is used above to simplify the discussion, the rewrite rules  
5 are actually presented in an extension of SQL. For example, the first example of the rewrite rule noted above might appear as:

REWRITE

10           SELECT eid, loc  
              FROM MAP  
              WHERE INSIDE(windows, loc)  
  
              AS  
              SELECT eid, loc  
              FROM MAPCLIP  
              WHERE MAPCLIP.WINDOW = WINDOW

15

Second Aspect of the Invention

The second aspect of the invention concerns the approach used to generate 22 alternative queries using the rewrite rules 16. In this section, it is assumed for simplicity sake that the queries have no inequality constraints.

20           It is important to note that a subexpression which is equivalent to the left hand side of the rewrite rule can be replaced by the right-hand side of the rewrite rule to derive a new query. Although the replacement step is easy, determining whether a subexpression is equivalent to the left hand side of a rewrite rule is more difficult. Thus, a novel procedure to check equivalence of two conjunctive  
25           queries is described below.

Using conventional knowledge, one might believe that two conjunctive queries are equivalent if and only if there is a renaming of variables such that there is a one-to-one mapping among literals in the query. In other words, two queries are equivalent if and only if they are isomorphic. However, merely substituting the right-side of a rewrite rule for a subexpression found to be  
30

equivalent to the left-side of the rewrite rule is not sufficient. The following example shows that such simple substitution is not enough to ensure equivalence.

Example 4:

The left-hand side of the following rewrite rule matches the query

5       $Q_1$ .

$\text{Business}(\text{name}, \text{"Restaurant"}, \text{earn}, \text{size}, \text{eid}), \text{Map}(\text{eid}, \text{loc}) \Leftrightarrow$   
 $\text{Map\_Restaurant}(\text{eid}, \text{loc})$

Hence, the rewrite rule can be used to generate alternative query

$Q'_1$  from query  $Q_1$ .

10      $Q_1(\text{loc}) :- \text{Business}(\text{bizname}, \text{"Restaurant"}, \text{earn}, \text{size}, \text{eid}),$   
 $\text{Map}(\text{eid}, \text{loc}), \text{Owner}(\text{bizname}, \text{"bob"})$   
 $Q'_1(\text{loc}) :- \text{Map\_Restaurant}(\text{eid}, \text{loc}), \text{Owner}(\text{bizname}, \text{"bob"})$

Note  $Q'_1$  is the effect of using the simple substitution approach to  $Q_1$ . Assume that bob owns only motels. The query  $Q_1$  returns an empty set, but the query  $Q'_1$  need not.

15     The crux of the problem with simple substitution is that the semantics of rewrite rules ensure equivalence of both sides of the rule only over universal variables. Therefore, it is not enough for a mapping to satisfy the condition of one-to-one mapping. The mapping has to satisfy additional constraints. Thus, a novel procedure to check equivalency of two conjunctive queries must satisfy the one-to-one mapping as well as the additional restraints. This novel procedure is termed substitution mapping.

20     Substitution mapping is defined as follows. Let  $I \Rightarrow r$  be a rewrite rule and  $Q$  be a query. A mapping from variables of  $I$  to that of  $Q$  is called a substitution mapping from the rule to the query if:

- (a) the mapping is one-to-one from literals in  $I$  to literals in  $Q$ ;
- (b) only universal variables in  $I$  are mapped to (if at all) constants in  $Q$ ;
- (c) images of the non-universal variables may not be among the projection variables of  $Q$ ;

- (d) the images of the non-universal variables of  $l$  may not appear in any literal of  $Q$  that is not in the image of  $l$ ; and
- (e) a constant in  $l$  can only map to itself.

In the definition, the term "image" of a variable (or a literal) with respect to a mapping means the variable (or the literal) to which the former is mapped. Likewise, the image of a query with respect to a mapping, refers to the literals obtained by applying the mapping to the query.

Comparing the above definition with the mapping performed in Example 4 shows that in this case the simple mapping performed violated condition (d) since  $bizname$  is an image of the non-universal variables. Consider a variant of the query in Example 4 where the literal  $Owner$  is replaced by  $Historic(loc)$  in  $Q$ . In such a case, the condition (d) is satisfied and the mapping satisfies substitution mapping.

In order to generate an alternative query, a rewrite rule is applied to a query using two basic steps. First, identify if there exists a subexpression such that there is a substitution mapping from the rule to the query. Second, substitute the subexpression with the right-hand side of the rule (after renaming).

A procedure to generate all equivalent conjunctive queries obtained by application of a rule to a query is as follows. All possible mappings which satisfy the condition of one-to-one mapping could be determined, and then the procedure could determine whether the mappings satisfy the additional constraints of substitution mapping. For queries with no repeated table names, applications of a rule can be performed in time  $O(n \log n)$  in the size of the query and the rewrite rule ( $n$ ). Since in practice most queries fall in this category, application of rules introduces little overhead for most queries.

The objective in introducing rewrite rules is to present the optimizer 11 with alternatives that arise due to semantics of foreign tables. The alternative queries are generated 22 by application(s) of the rules. Each of the generated queries is equivalent to the given query.

The closure of the query represents the set of all the alternatives that can be generated by applications of rewrite rules. More particularly, the closure of a query Q with respect to a set R of rewrite rules is the set of queries:

$$\text{closure}(R, Q) = \{Q' \mid Q \Rightarrow_R^* Q'\}$$

5 The symbol  $Q \Rightarrow_R^* Q'$  is used to denote the fact that  $Q'$  has been obtained from Q by a finite sequence of applications of rewrite rules (drawn from the set R).

The rewrite rules can also be used in a sequence as illustrated by the following example.

Example 5:

10 Consider the following query.

$Q(\text{loc}) :- \text{Map}(\text{eid}, \text{loc}), \text{Inside}(\text{w1}, \text{loc}), \text{Inside}(\text{w2}, \text{loc})$

We can apply the rule in third example of the rewrite rules to generate the query  $Q'$ :

$Q'(\text{loc}) :- \text{Map}(\text{eid}, \text{loc}), \text{Inside}(\text{w}, \text{loc}), \text{Intersect}(\{\text{w1}\}, \{\text{w2}\}, \text{w})$

15 Thereafter, an application of the first example of the rewrite rule results in the following query  $Q''$ .

$Q''(\text{loc}) :- \text{Mapclip}(\text{eid}, \text{loc}, \text{w}), \text{Intersect}(\{\text{w1}\}, \{\text{w2}\}, \text{w})$

20 Traditionally, optimization is concerned with choosing an optimal plan for a query. However, in the presence of rewrite rules, the invention causes a set of equivalent queries (closure) to be produced. Thus, the optimization problem is to pick the cheapest among the optimal plans of the queries in the closure of the given query. The optimization approach of the invention is to first generate 22 a set of equivalent queries (Step 1), and then choose 24 the cheapest among the optimal plans for each query obtained (Step 2). Step 1 is detailed below, and Step 25 2 is described in the third aspect of the invention.

Given a query and a set of rewrite rules, a closure procedure computes the closure of the query with respect to the set of rewrite rules. The closure procedure is illustrated in Figure 3.

30 Figure 3 is a flowchart of an embodiment of the closure procedure according to the second aspect of the invention. Closure is the set of equivalent queries produced from the input query and the rewrite rules. The procedure

begins by deciding 31 whether a new query is available for processing. On the first iteration, the new query is the input query 10 and it is available. On subsequent iterations, the new queries are the alternative queries which have been generated by other portions of the procedure.

5           In any case, if there are no queries available, the procedure is completed. On the other hand, if it is determined 31 that one or more new queries are available, then a next new query is obtained 33. Next, a decision 34 is made based on whether at least one rewrite rule is available to be applied. This decision 34, at minimum, would determine if, for example, a invoking a foreign function  
10          had at least one rewrite rule associated therewith. If no rewrite rules are available for use with the query, then the procedure returns to decision 31 to process the next query.

15         On the other hand, if it is determined 34 that at least one rewrite rule exists, then a next rewrite rule is obtained 35. The next rewrite rule is one of the rewrite rules which are available for use with the query. The next rewrite rule is then used in a decision 36 which determines whether a substitution mapping can be performed from the rewrite rule to the query. The requirements of substitution mapping were discussed in detail above. If the substitution mapping fails, the processing returns to decision 34 to determine if another rewrite rule is available.  
20         If the substitution mapping is successful, a new query is generated from the rewrite rule (in use) and the query (in use). Thereafter, the processing loops back to decision 36 for cases where the rewrite rule has more than one substitution mapping for the query in use (not the newly generated queue).

25         An example of programming code for a particular closure algorithm (gen\_closure) is contained in Appendix A. The algorithm repeatedly invokes the function rewrite which is also contained in Appendix A. Given a rewrite rule r and a query Q, the function rewrite(r,Q) derives all queries that can be obtained by a single application of r to Q. In Appendix A, the phrase "sound application" refers to successful substitution mapping. The algorithm gen\_closure is iterative and during each iteration, there is a set of queries (new) which acts as the seeds to generate additional queries by invoking rewrite. Only those derived queries that  
30

were never generated before, act as seeds to the next iteration. The algorithm gen\_closure is a sound and complete procedure to obtain closure.

In order to access the relevant rules efficiently, it is preferable to maintain the rules in a rule-table (see Figure 4) which is indexed on the conjuncts that appear on the left-hand side of the rule. It may also be advantageous to use rule classes which allow the rules to be partitioned into classes and then applied such that all rules in a given class would be applied before any rules in another class.

Generally speaking, in the closure procedure shown in Figure 3 termination depends on whether the closure of the query with respect to a set of rewrite rules is finite or not. Thus, the procedure terminates if and only if the closure is finite. However, there could exist queries and rewrite rules such that the closure of a query with respect to a set of rewrite rules is an infinite set. As a result, several conditions are identified to facilitate closure.

Two conditions are described to ensure closure if a set of rewrite rules satisfies either of the conditions. The first condition is a Non-Increasing Length Property condition. This condition requires that the length of the right-hand side of every rule be shorter or equal to its left-hand side. For example, a domain-calculus expression length may be defined as the number of conjuncts in the expression. Thus, a rewrite rule satisfies the non-increasing length condition if the right-hand side of the rewrite rule is no longer than its left-hand side. As an example, the following rewrite rule fails to satisfy the non-increasing length property condition.

Map\_Restaurant(eid, loc)  $\Rightarrow$  Business(name, "Restaurant", earn, eid),  
 Map(eid, loc)

The second condition is a Non-Null Property condition. The non-increasing length property fails to capture (i.e., over inclusive) situations where although there is a local violation of non-increasing length property condition, the length of the rewritten query can not increase beyond a certain limit. The non-null property condition captures these situations. This test for termination of a set of rewrite rules R has two steps. First, a set of inequalities Ineq(R) is derived from a set R of rewrite rules. The set Ineq(R) of inequalities is constructed as follows:

- Associate an integer variable with every table name.
- Obtain an algebraic expression for every domain calculus expression by: replacing every literal by the variable corresponding to the table name, and replacing conjunction by addition.

5        For every rewrite rule of the form  $\text{lhs} \Rightarrow \text{rhs}$ , an arithmetic inequality  $\text{left\_Exp} \geq \text{Right\_Exp}$  is formed, where  $\text{left\_Exp}$  and  $\text{right\_Exp}$  are the algebraic expressions for  $\text{lhs}$  and  $\text{rhs}$ , respectively. In addition, every variable is greater or equal to zero. The set of arithmetic inequalities corresponding to a set  $R$  of rewrite rules will be denoted by  $\text{Ineq}(R)$ . A set of rewrite rules  $R$  has finite closure for every query if no variable in  $\text{Ineq}(R)$  is null. A variable in  $\text{Ineq}(R)$  is said to be null if it must be assigned zero for every solution to  $\text{Ineq}(R)$ . Efficient algorithms based on linear programming techniques are known and available to check whether a variable in a set of linear inequalities is zero in all solutions.

10      The non-null property condition is demonstrated by the following example.  
15      Consider the set of cyclic rewrite rules involving Map\_Restaurant in the second example of the rewrite rules which fails to satisfy the non-increasing length property. The corresponding inequalities are ( $m$  for Map,  $r$  for Restaurant,  $d$  for Map\_Restaurant):

$$\{m \geq 0, d \geq 0, r \geq 0, d \geq m+r, m+r \geq d\}$$

20      Evidently, there exists no variable which is null and therefore there is finite closure.

25      In addition to above conditions, heuristics can also be used to ensure termination of the closure procedure when neither of the termination conditions (e.g., non-increasing length property and non-null property) are satisfied. In such a case, the closure procedure enumerates only a finite subset of closure. Techniques to specify such a subset are discussed below. Furthermore, if desired, heuristics can also be used to eliminate some queries generated in Step 1 from consideration in Step 2 of the optimization process.

30      The closure of a query is typically limited to few queries. Nonetheless, heuristic techniques are useful in restricting the number of candidate equivalent queries that will be considered for optimization in the second phase of the

optimization. These heuristic techniques can be of two types. First, only a subset of the closure can be enumerated. Next, some of the alternative queries generated may be eliminated from consideration in Step 2 of the optimization process.

For selective enumeration of closure, one can use a budget to determine the maximum time spent on enumeration. The closure procedure can be easily extended to incorporate this modification. Another alternative is to modify the rewrite rules. This alternative approach is illustrated by the following example.

Example 6:

A rewrite rule such as:

Business(name, "Restaurant", earn, size, eid), Mapclip(eid, loc, window)  $\Rightarrow$  Business(name, "Restaurant", earn, size, eid),  
 Mapclip(eid, loc, small\_win), Intersect(window, w, small\_win)

can be modified to be as follows:

Business(name, "Restaurant", earn, size, eid), Mapclip(eid, loc, window)  $\Rightarrow$  Business(name, "Restaurant", earn, size, eid),  
 SpecialMapclip(eid, loc, window)

The modified rewrite rule contains a new table name SpecialMapclip.

Unlike the original (unmodified) rule, the modified rewrite rule can not be used repeatedly since there is no rewrite rule where SpecialMapclip occurs in the left-hand side. However, after closure is generated, the expression "Mapclip(eid, loc, small\_win), Intersect(window, w, small\_win)" will be substituted for "SpecialMapclip(eid, loc, window)".

Hence, the example shows that a subexpression can be treated as a single literal, thereby reducing the size of the effective closure. The choice of subexpression is used to decide the subset of closure that is selected for enumeration. By using this strategy, one can ensure that any set of rewrite rules satisfy the non-length increasing property condition and thereby ensure termination.

Another application of the heuristics is in identifying equivalent queries that need not be optimized. This may be determined by a crude cost estimate or by

using the following approach. Some of the rewrite rules (e.g., the third example of the rewrite rules) may be marked as "always improving". If such a rule derives a query  $Q'$  from a query  $Q$ , then the query  $Q$  need no longer be optimized since it is assumed that  $Q'$  will always result in a better optimal plan.

5      Third Aspect of the Invention

Query optimization in the presence of rewrite rules and foreign tables require new dimensions to the traditional optimization problem. First, the presence of foreign tables require introduction of new join methods as well as cost models that are appropriate for foreign tables. In this section, it is assumed 10 that the cost model assigns a real number to any given plan in the execution space, and satisfies the principle of optimality (see Cormen et al., *Introduction to Algorithms*, The MIT Press, 1990), which is implicit in all relational optimizers that use dynamic programming. Further, since the foreign tables may have safety constraints, the traditional join enumeration phase must ensure that only those 15 reordering of the joins which satisfy the safety constraints are considered. In other words, one needs to ensure that the bindings that are passed to the foreign function satisfy the safety constraints. This is a well-studied problem. See, e.g., Ullman, *Principals of Database and Knowledge-Base Systems*, Computer Science Press, 1989 for a discussion. Finally, the task is to choose 24 an optimal plan from an 20 enriched space, which contains multiple equivalent queries.

Of course, any discussion of optimality must be with respect to a space of possibilities. The execution of a query is traditionally represented syntactically as annotated join trees (see Krishnamurthy et al., *Optimization of nonrecursive queries*, Proceedings of International Conference on Very Large Data Bases, pp. 25 128-137, Kyoto, Japan, August 1986) where the internal node is a join operation and each leaf node is a base table. In accordance with the invention, a leaf node can also be a foreign table. The optimizer considers all left-deep annotated trees, thus forcing a linear ordering of joins. Therefore, the execution space is defined to be the space of all left-deep join trees for each equivalent query obtained from Step 1 of the optimization process. Step 2 of the optimization process is then to 30 choose a plan of least cost from the execution space.

5

Figure 4 is a block diagram of an embodiment of the optimization process according to the invention. Figure 4 illustrates the functional units associated with the optimizer 11 shown in Figure 1. In particular, the optimization process makes use of an equivalent query generator 40 which uses the rewrite rules 12 stored in a rule table 42 to produce the equivalent queries. The optimization process also makes use of a reorderable unit generator 44 and a query optimizer 46. The query optimizer 46 is able to operate efficiently by interacting with local optimizers 47 and a plan table 48. The optimization process according to the invention is discussed in more detail below.

10

Step 2 of the optimization process is performed by the query optimizer 46 which makes use of a cost optimization algorithm to determine or select the optimal plan from the execution space. Two different approaches for the cost optimization algorithm are described below.

15

20

The first approach for the cost optimization algorithm is a direct extension of the traditional dynamic programming optimization algorithm. Since the execution space is the union of the execution spaces of the equivalent queries, the optimal plan for each query in the closure can be obtained and then selection of the plan that has the least cost among them. For example, if  $Q_1$  and  $Q_2$  are the only two queries in the closure, one can determine the optimal plans  $P_1$  and  $P_2$  for  $Q_1$  and  $Q_2$  respectively, and then select the cheaper plan of  $P_1$  and  $P_2$ . Therefore, the following cost optimization algorithm is a direct extension of the traditional dynamic programming algorithm and includes the steps:

25

(i) optimize each query using the traditional dynamic programming algorithm and obtain the best plan for the query.

(ii) choose the cheapest among the best plans obtained in Step (i).

30

Observe that step (i) can be accomplished by any traditional relational optimizer (see, e.g., Selinger et al., *Access path selection in a relational database management system*, Proceedings of the ACM SIGMOD International Symposium on Management of Data, pp. 23-34, Boston, June 1979) with the straight forward extension to check for safety while considering a join order. The space requirement for this algorithm is the maximum space required for optimization of

any of the equivalent queries. However, the algorithm has a poor time complexity since it fails to take advantage of the fact that there is commonality among queries in the closure.

The second approach for the cost optimization algorithm is preferred  
 5 because it operates much faster than does the direct extension. The key to this approach is that the cost optimization algorithm exploits commonality among queries across multiple equivalent queries during optimization. Although the traditional dynamic programming based algorithm exploits the commonality among the subqueries of a single query to reduce the time complexity of the optimization,  
 10 it does not exploit the commonality among subqueries across multiple equivalent queries during optimization.

Before describing the improved cost optimization algorithm, a description of how the subqueries that are shared by two or more equivalent queries can be identified. This identification process is performed during Step 1 of the  
 15 optimization procedure where equivalent queries in the closure are generated 22 by applications of rules. The example which follows illustrates the identification process.

Example 7:

Consider an application of the second example of the rewrite rule to  
 20 query Q that results in the query Q'. The MapEngine file *Historic* contains the location of all historic sites and the foreign table *price*<sup>b</sup>(loc, amount) provides the value of the real estate. The following query asks for prices of real estate of the places which are historic and have restaurants.

25           Q(amount) :- Historic(loc), Business(bizname, "Restaurant", earn, eid), Map(eid, loc), Price(loc, amount)  
               Q'(amount) :- Historic(loc), Map\_Restaurant(eid, loc), Price(loc, amount)

Observe that the subexpression "Historic(loc), Price(loc, amount)" is common in  
 30 both the queries. This commonality can be detected at the time the rewrite rule is applied by observing that the first and the last literals are left unchanged by the

application of the rewrite rule. Since as part of optimizing the query  $Q$ , the optimal plan for the common subquery  $\text{Historic(loc)}$ ,  $\text{Price(loc, amount)}$  must be determined, it can be reused when determining the optimal plan for the query  $Q'$ .

In accordance with the improved cost optimization algorithm, each query is  
 5 optimized one at a time, but for any subquery that is shared with another equivalent query (only those subqueries), its optimal plan is stored in a plan table 48 so that its optimal plan does not have to be rederived. Thus, when constructing the optimal plans, the plan table 48 is consulted to check whether the plan already exists.

An example of the programming code for the improved cost optimization algorithm is provided in Appendix A (i.e., procedure Optplan). Unlike the traditional algorithm, the improved algorithm follows a top-down approach. The base cases in Optplan where the query has only one or only two conjuncts have been omitted. The improved algorithm invokes local optimizers to generate  $Q_i$   
 10 from  $S_i$  and  $q_i$  (see Appendix A). The local optimizers 47 make use of information about the cost-model. According to the improved cost optimization algorithm optimal plans are built bottom-up, i.e., the candidate optimal plans for a query are constructed from optimal plans of its immediate subqueries. As in traditional optimizers, the optimizer 11 treats the conjuncts (e.g.,  $\text{salary} > 50k$ )  
 15 that are built-in boolean conditions (sargable predicates) specially. They are not reordered but instead their effect is considered for selectivity factors as well as for determining access paths in the local optimizer. Finally, note that the improved algorithm discards the join orders which do not satisfy the safety constraints (not shown in Appendix A).

When accessing the plan table 48, it is beneficial to ensure that searching  
 20 for the optimal plan of a query is efficient. One efficient approach is to maintain the plan table as a hash-table where each query is represented as a sorted list of integers. Every integer corresponds to the index of the entry a global literal table where literals of the queries are stored. Therefore, comparison of two queries  
 25 amounts to checking whether two strings are equal and thus can be done very

efficiently. The following example illustrates the representation of the literals as well as how the commonality among queries are captured in this representation.

Example 8:

This example makes use of the queries  $Q$  and  $Q'$  of Example 7.

5 Assume that the indexes for the global literal table corresponding to the literals in the query  $Q$  are 1,2,3, and 4 from left to right.

Therefore,  $Q$  will be represented by the string (1234). However, once the rewrite rule is applied, a new literal

10 Map\_Restaurant(eid,loc) is created which is entered with an index 5 in the global literal table. The representation for  $Q'$  will be (145).

The optimization of the query  $Q$  creates the optimal plan for (14) which is then stored in the plan table. During the optimization of the query  $Q'$ , first the plan table is consulted to see whether a plan for (145) already exists. Since it does not exist, the optimal plans 15 for each subplan are recursively constructed. In particular, before constructing the optimal plan for (14), the plan table is consulted and the optimal plan of (14) is reused for optimization.

20 The improved cost optimization algorithm has the desirable features that for no subquery, a subplan is rederived. Moreover, only plans for shared subqueries are retained in the plan table.

It is well-known that dynamic programming based 25 algorithms can be presented either as top-down or bottom-up (see Cormen et al., *supra*). Therefore, a variant of the cost optimization algorithm can be created, namely a bottom-up approach. Indeed, there can be two possible variants in a bottom-up approach. One possibility is to optimize all the equivalent queries together. Thus, optimal plans for all subqueries of size  $n$  are constructed before any optimal plan for any subquery of size  $(n+1)$  is constructed. This approach has the advantage that it requires less space than the top-down approach since to construct any plan of size  $(n+1)$ , only plans of size  $n$  are needed and so space 30 may be reused. Unfortunately, since the subqueries for all equivalent queries are constructed together, the time for the completion of the optimal plan for any of the

equivalent queries is longer than the top-down approach. Another variant of the bottom-up approach is where optimization is done one query at a time but the optimal plans of shared subqueries are saved. While this rectifies the shortcoming of the previous approach, it suffers from the problem of not being able to share the largest common subquery as the following example illustrates.

5

Example 9:

10

Assume that there are two queries in the closure with the representations (1234) and (1235). The queries share the subquery (123). Assume that the optimal plan for (1234), and hence for (123), have already been constructed. Consider the task of building an optimal plan for (1235). If the optimizer uses the bottom-up approach, it fails to recognize the existence of the cached plan (123) until all smaller plans are constructed and therefore, unnecessary accesses to plan table will be made.

15

Thus, there are pros and cons of using a bottom-up or a top-down version of the optimization algorithm. It may be preferable to use the top-down variant augmented with a branch and bound strategy. Thus, if a partial plan is found to have exceeded the cost of the optimal plan that has been found so far, then that partial plan need not be completed since it is guaranteed to be suboptimal.

20

Heuristic techniques may also be incorporated into the optimization process. For example, heuristics may be used to determine the order in which queries are optimized. Also, it may be beneficial to discard search for an optimal plan (instead of enumerating the entire space) after a significant improvement over the optimal plan for the given query is obtained.

25

Inexpensive and affiliated predicates may also play an important role in optimization. In a traditional relational optimizer, selection conditions are not reordered like other database relations. Rather, the selection conditions are evaluated as early as possible. Indeed, the cost of reordering joins is exponential in the number of literals being reordered. Therefore, unnecessary reordering introduces overhead for the optimizer. Consider a foreign table such as Inside(w, loc) which checks whether a point loc is inside window w. The invocation of such

a table is inexpensive. Therefore, optimization costs can be decreased by designating certain foreign tables as inexpensive tables. The literals in a query that correspond to inexpensive tables are not reordered, while the rest of the literals are reordered. The literals corresponding to inexpensive tables are 5 evaluated as early as possible in safe ordering of reorderable literals. The presence of inexpensive conjuncts introduce the step of generating reorderable units by the reorderable unit generator 44 of Figure 4. The sargable predicates in a relational optimizer have the property that they may be pushed down to the storage system layer. See Selinger et al., *supra*. Likewise, the facility to indicate 10 which inexpensive tables are affiliated to a reorderable table may also be provided. This information may be exploited during optimization.

The query processing techniques for foreign functions used by the 15 optimizer 46 are now described. In particular, the focus is on query processing techniques for the join operation. In the subsequent description of the cost model, it is assumed that the foreign table occurs as the right child of a join node in a left-deep tree. Therefore, in a left-deep join tree, it is assumed that the foreign table joins with some other table (intermediate or base), which is referred to as left table.

The key point is that any access to a foreign table must respect the safety 20 constraints. Therefore, before a tuple can be obtained from a foreign table, it must be passed the bindings that are required due to safety constraints. In most systems, for each binding, an invocation is made for the foreign function. Such a technique is inefficient for invoking foreign functions which have high cost of invocation and return many tuples as output. Therefore, according to the 25 invention, query processing for foreign functions consists of viewing the "join" with a foreign table to have two phases: invocation and residual join. The invocation phase consists of passing the values for the bound arguments of the foreign table from the left table. There are several ways in which this invocation may occur:

Simple Invocation: In this scheme, for each tuple in the left table an invocation is made for the corresponding bound arguments, as specified in the safety constraint.

5

Group Invocation: In this scheme, for each distinct values of the bound arguments from the left relation, a single invocation is made for the bound arguments.

10

The group invocation technique adds the overhead of identifying the set of distinct values for the bound arguments. However, it has the advantage of fewer invocations, which is important for foreign tables for which each invocation is expensive (for example, computing the mean distance of a point from all other points). Moreover, if the left table is already sorted on the bound arguments prior to join, then group invocation is superior.

15

Since an invocation generates a set of tuples, the step of residual join is similar to a traditional join and any join method may be used. Also, the selection conditions that apply to one or more free (f) arguments of the foreign table are evaluated during this phase. The simplest choice for the residual join is nested loop where the tuples generated for each invocation are treated as the inner table for the join. The only difference is that the inner table changes (potentially) for each invocation. This residual join method can be combined with the two techniques for invocation.

20

The combination of simple invocation and the choice of nested loop join technique results in a join algorithm which is similar to the traditional nested loop join and is referred to as foreign nested loop join (FNL). The combination of group invocation and the choice of nested loop join results in an algorithm similar to the sort merge join and is referred as foreign sort-merge join (FSM). An example of programming code for the FSM algorithm appears in Appendix B (i.e., Foreign Sort Merge Join Algorithm). The FSM algorithm is preferred when the invocation of foreign tables is expensive. Finally, to reduce the number of invocations, catching results of invocation was suggested in the context of Postgres and such an alternative can be used with our approach to query processing. See, e.g., Stonebraker et al., *supra*.

25

30

The cost model must be able to compute the cost of any given plan.

For traditional relational optimizers, a descriptor for a table includes statistical information about the table such as the number of unique values in each argument position (i.e., in each column) and the expected number of tuples in the table.

5 The cost model uses the descriptors to compute the cost of an operation (e.g., a join). The cost model also produces a new descriptor which contains the statistical information of the intermediate table which is obtained after the operation.

Although the approach of the invention to the cost model is to preserve the relational descriptor for database and intermediate tables, the invention requires 10 two extensions. First, a descriptor for foreign tables must be provided. Second, explanations of on how such a descriptor can be used to obtain the cost of an operation and how one can derive the relational descriptor for the intermediate table after a "join" with the foreign table.

For each foreign table, the following information can be registered:

15 Safety Constraints: This information is not directly used by the cost model, but is used by the optimizer.

Cost: The cost of invoking the foreign table once.

Fanout: The number of "output tuples" expected for each invocation.

20 For each attribute:

Domain Size: The size of the representation of each domain element. The cardinality of the domain. A permissible assignment to cardinality is infinite.

25 Unique Value Factor: The expected number of unique values the attribute has for each invocation. If this parameter is not explicitly provided, the fanout is used to approximate this factor. If all the domains are finite, uniform distribution assumption is used to compute this factor as well.

One should observe that except for safety, all the other properties of the 30 descriptor are relevant for the cost model. The cost model is an extension of the model proposed in Chimenti et al., *Towards an open architecture for LDL*,

Proceedings of the 15th International VLDB Conference, pp. 195-203,  
 Amsterdam, August 1989. One should also observe that the parameters in the  
 descriptor need not necessarily be constants, but can depend as well on any  
 constants that appear in the query at compile-time. An example of a descriptor for  
 5 a foreign is provided below.

Example 10:

The descriptor for the foreign table  $\text{Intersect}^{\text{bbf}}(\text{window1}, \text{window2},$   
 $\text{window3})$  could be characterized by a cost of .012ms, a fanout of 1,  
 10 unique value factor of 1. The size of each domain element is that  
 corresponding to a real and has domain size - infinite}. Observe  
 that fanout and unique value factor is 1.

To compute the descriptor extensions are needed. For simplicity, only the  
 situation when the foreign table occurs as a right child of left deep join trees is  
 described. Since it is assumed that the foreign table appears as some right leaf  
 15 node, one can also assume the existence of a descriptor for the intermediate  
 relation, which is referred to as the left relation. In the optimizer 46, one can  
 register a customized function to compute the resulting descriptor. Such a  
 function can take as its argument the descriptor for the left relation. Below is  
 described a default way to compute the descriptor for the intermediate table.

20 Before, describing the formulas to derive the descriptor, it is useful to  
 explain the left uniqueness factor. The left uniqueness factor estimates the  
 expected number of distinct invocations of the foreign table for the given  
 descriptor for the left relation. There exist several ways to approximate the left  
 uniqueness factor. One way is to use a simple formula  
 25 for the left uniqueness factor as described below.

For the foreign table, one or more argument positions will be required to  
 be bound. Therefore, there exists a corresponding set of attributes A in the left  
 relation which provide the values for the bound arguments of the foreign table.  
 Let P be the product of the expected number of unique values for the set of  
 30 attributes A in the left relation. The descriptor of the left relation is used to  
 compute P. Clearly, the number of distinct invocations can not exceed P.

However, the number of distinct invocations also may not exceed  $N$ , the number of tuples in the left relation. Therefore,  $\min(P, N)$  can be used as the left uniqueness factor. Observe that our formula provides an upper bound of the left uniqueness factor. An example shows an application of the formula.

5

Example 11:

Consider the following query which provides the location of the terminals for the bus routes.

Query(route, loc) :- Terminal(route, eid), Map(eid, loc)

10

Assume that the descriptor for Terminal has 100 tuples

but the number of expected unique values in the second argument is expected to be 10. In such a case, the left uniqueness factor will be estimated as 10.

15

In the following description, it is assumed that there are no selection conditions other than equality between the left table and the bound arguments of the foreign table. The effect of output selection conditions as well as effect of projection on foreign tables on the descriptor are computed by treating the result of the join of left relation with the foreign table as an intermediate table (like any interior node of the join tree). Therefore, the cost formulas provided below are for the invocation phase only.

20

Number of Tuples: The estimated number of tuples following the join is  $N' = F*N$ , where  $F$  is the fanout of the foreign table and  $N$  is the number of tuples in the left relation.

25

Number of Unique Values: The estimated number of unique values corresponding to the  $i^{\text{th}}$  argument of the foreign table is given by:

$UVF_i * UI$  where  $UVF_i$  is the unique value factor for the  $i^{\text{th}}$  attribute.

The parameter  $UI$  is the left uniqueness factor, discussed earlier.

30

Cost: Cost of foreign nested loop and foreign sort-merge join are provided. Assume that  $N$  is the number of tuples in the left relation,  $C$  is the cost of invoking the foreign table and  $UI$  is the uniqueness factor. The following costs are for the invocation phase only.

- Foreign Nested Loop: C\*N

- Foreign Sort-Merge: Cost<sub>sort</sub>(N) + UI\*C

The many features and advantages of the present invention  
are apparent from the written description and thus it is intended by  
5 the appended claims to cover all such features and advantages of the  
invention. Further, since numerous modifications and changes will  
readily occur to those skilled in the art, it is not desired to limit the  
invention to the exact construction and operation as illustrated and  
described. Hence, all suitable modifications and equivalents may be  
10 resorted to as falling within the scope of the invention.

CLAIMS

1. A relational database system, comprising:  
a relational database (15) having a plurality of tuples;  
at least one foreign function (16) having declarative rewrite rules (20)  
associated therewith; and  
an optimizer (11) for optimizing a query accessing said relational database  
and said at least one foreign function using cost information.
2. A relational database system as recited in claim 1, wherein said rewrite  
rules are used to generate (22) equivalent queries from which an optimal plan is  
selected (24).
3. A relational database system as recited in claims 1 or 2, wherein said  
rewrite rules (20) associated with each of said foreign functions are expressed in a  
high-level declarative language.
4. A relational database system as recited in claims 1, 2 or 3,  
wherein the high-level declarative language used for said rewrite rules is an  
extension of SQL having the form REWRITE Query 1 AS Query 2, and  
wherein the rewrite rules are stored in a rule table (42) indexed on the  
conjuncts that appear on the left-side of the rewrite rules.
5. A relational database system as recited in claims 1, 2, 3 or 4, wherein said  
optimizer (11) exploits commonality among queries across multiple queries, and  
wherein said relational database system further comprises a plan table (48)  
which stores previously determined optimal plans for subqueries to exploit the  
commonality.
6. A relational database system as recited in claims 1, 2, 3, 4 or 5, wherein  
the cost information comprises cost information on the foreign functions.

7. A method for optimization of a query (10) invoking relational database tables (15) and foreign functions (16), said method comprising:
  - (a) providing (20) rewrite rules for the foreign functions;
  - (b) receiving an input query to be executed;
  - 5 (c) generating (22) alternative queries from the input query and the rewrite rules; and
  - (d) generating (24) an optimal plan.
8. A method as recited in claim 7, wherein said generating step (c) comprises determining if a left-side of at least one of the rewrite rules is equivalent to at least a portion of the query using substitution mapping.
9. A method as recited in claims 6 or 7, wherein said generating step (d) comprises:
  - (d1) generating optimal plans for the input query and for at least one of the alternative queries; and
  - (d2) selecting one of the optimal plans.
10. A method as recited in claim 9, wherein said generating step (d1) exploits commonality among queries across multiple queries.
11. A relational database system substantially as herein described with reference to the accompanying drawings.
12. A method substantially as herein described with reference to the accompanying drawings.

APPENDIX A

```

Function gen_closure( $R, Q$ )
begin
   $S = Q$ ;
   $\delta = Q$ ;
  repeat
    new =  $\emptyset$ 
    for each  $q$  in  $\delta$  and  $r$  in  $R$  do
      new = new  $\cup$  rewrite( $r, q$ );
    endfor
    if new  $\subseteq S$  then return( $S$ );
     $\delta = \text{new} - S$ ;
     $S = S \cup \delta$ ;
  forever
end

```

```

Function rewrite( $r, Q$ )
begin
   $Q_r = \emptyset$ 
  for every sound application  $A$ 
  of  $r$  to  $Q$  do
     $Q_r = Q_r \cup \{A_Q\}$ 
    where  $A_Q$  is the derived query
    due to application  $A$ 
  endfor
  return( $Q_r$ )
end

```

Algorithm to Compute Closure

```

Procedure OptPlan( $Q$ ) :
begin
  if exists optimal( $Q$ ) then return;
  Let  $Q = (q_1, \dots, q_n)$ ;
  Let  $S_i = Q - \{q_i\}$ ;
  for each  $i$  do
    OptPlan( $S_i$ );
     $Q_i = \text{Plan for } Q \text{ from } S_i \text{ and } q_i$ 
  endfor;
  Choose best among  $Q_i$ 
  and add to plan table.
end

```

Optimization Algorithm

APPENDIX B

Function  $FSM(Left, FTable)$   
 (Left is left relation, FTable is a Foreign Table)  
 begin  
    $Join = \emptyset$   
    $Temp\_Left = GROUPBY(Left, Bound)$   
   where sorting and Grouping is on bound arguments of FTable  
   for every group  $L_i$  of  $Temp\_Left$  do  
      $FT_i = Invoke(FTable, Bval_i)$   
     where  $Bval_i$  are the values in partition  $L_i$   
     for bound arguments of  $Bound$   
      $Join = Bag\_Union(Join, Merge(L_i, FT_i))$   
   endfor  
   return( $Join$ )  
 end

**Foreign Sort Merge Join Algorithm**